KAIS Preprint

Discriminative Learning of the Prototype Set for Nearest Neighbor Classification

Shin Ando

arXiv:1509.08102v5 [cs.LG] 21 Aug 2016

Abstract The nearest neighbor rule is one of the most widely used models for classification, and selecting a compact set of prototype instances is a primary challenges for its applications. Many existing approaches for prototype selection exploit instance-based analyses and locally-defined criteria on the class distribution, which are intractable for numerical optimization techniques. In this paper, we explore a parametric framework with an adjusted nearest neighbor rule, in which the selection of the neighboring prototypes is modified by their respective parameters. The framework allows us to formulate a minimization problem of the violation of the adjusted nearest neighbor rule over the training set with regards to numerical parameters. We show that the problem reduces to a large-margin principled learning and demonstrate its advantage by empirical comparisons with recent *state-of-the-art* methods using public benchmark data.

Keywords Nearest neighbor classifier, Prototype selection, Soft maximum, Large-margin learning

1 Introduction

Applying the nearest neighbor rule based on a set of *prototype* instances is one of the most widely used models for classification and a typical example of non-parametric, instance-based learning algorithms [31,40]. *Prototype selection* is a supervised learning technique related to the nearest neighbor algorithm, which selects a set of prototypes from the labeled training data to store and to use in nearest neighbor prediction on the test data. The *prototypes* are selected by various strategies in order to achieve the goal of improving both computational efficiency and robustness of the nearest neighbor algorithm. It is a critical problem for practical applications due to the large cost of storing and computing distances in the neighbor algorithms [7].

There has been a substantial amount of literature on prototype selection published over the years in the topic of data mining. Many existing techniques for prototype selection have been motivated by instance-based or locally-defined criteria, e.g., removing prototypes far from the decision boundary of two classes to reduce redundancy, or conversely, removing those near the boundary for generalization. Improvements by sophisticated hybridization of such techniques have also been reported [37]. Due to the instance-based nature of these techniques, however, it is generally difficult to learn their optimal parameter values. Alternatively, critical parameters are selected empirically through cross validation and *wrapper* methods [31].

S. Ando
School of Management, Tokyo University of Science,
1-11-2 Fujimi, Chiyoda-ku, Tokyo, 102-0071, Japan
E-mail: ando@rs.tus.ac.jp

Recently, the emergence of *Big Data* has ignited more approaches to efficiently compute the nearest neighbor algorithm, such as Locality Sensitive Hashing and Approximate Nearest Neighbor Search [2,32]. ANNS methods yield excellent run-time performances without removing any of the training data from the prototypes. There are, however, a number of motivations for revisiting the topic of prototype selection. First, the state-of-the-art prototype selection methods have the advantage in terms of classification accuracy and are competitive in terms of efficiency compared to ANNS [18]. Secondly, ANNS is subject the constraints for the LSH function family, which require the use of standard distance functions and limit its applicable domains. Finally, a prototype selection method can also achieve the function of data summarization, which becomes more important for knowledge discovery with the emergence of larger datasets.

In this paper, we revisit the prototype selection problem to explore a parametric approach based on a discriminative learning principle. The key intuition of the proposed approach is two-fold. First, we define a parametric extension of the nearest neighbor rule, in which the choice of the nearest neighbor is adjusted by a numerical parameter assigned to each prototype. The parameter values are also used to infer whether the instances that are relevant or irrelevant for predicting over the training set. Secondly, we introduce a differentiable approximation of the condition that the nearest neighbor rule predicts the class of a training instance correctly, and subsequently derive a constrained optimization problem for minimizing the violation of the condition with regards to the numerical parameters of the prototypes. We further show that this problem reduces to a large-margin principled learning using a sparse representation of the relations between the neighboring instances.

The proposed method is also applicable in domains where only pairwise distances among the instances are available, while many techniques for prototype selection rely on the analyses of the feature vector space. These analyses are infeasible for non-vector input data and variables with complex dependency among them. The nearest neighbor algorithms are often used for such data, e.g., in time series classification, the Euclidean distance is often problematic due to the sequential structure of the time series [22], and a dissimilarity function based on non-linear warping is known to be highly effective [13,41].

In summary, the contribution of our work is a prototype selection approach which (a) enables an optimization algorithm directly addressing the violation of the nearest neighbor rule and (b) expands the application to the distance space input data. The rest of this paper is organized as follows. Section 2 discusses the related work. Section 3 describes the proposed method Section 5 shows the results of our empirical study. Section 6 presents our conclusion.

2 Related Work

2.1 Prototype Selection

The weaknesses of the nearest neighbor classifier in the large memory requirements, the time requirements for computing the similarities/dissimilarities, and the sensitivity to noisy labels has been well-known. Ohe of the ways to address these weaknesses is to reduce the number of prototypes at the pre-processing phase. The problem has been referred to as prototype selection [18], data reduction [39], instance selection [19,20,27], and template reduction [14,33]. The goal of the problem is to find a compact subset of the prototypes which maintains or possibly improve the generalization error of the nearest neighbor classification while decreasing the computational burden.

Some operations are commonly used in many prototype selection methods, e.g., an operation referred to as *editing*, eliminates or relocates *noisy* prototypes that cause misclassifications near the borders of different classes [7,39], and have the effects of smoothing and generalizing the decision boundaries. *Condensation* is another typical operation which discards prototypes far from the class borders in order to reduce the redundancy of the prototype set without affecting the decision boundary [3]. Combinations of these techniques and others such

as feature analysis [34] and clustering [30,37]. are used to analyze boundaries and interior instances in hybrid methods [26,36].

The implementations of these techniques are generally categorized into *wrapper* and *filtering* methods [31]. In the wrapper methods, different sets of prototypes are evaluated and selected based on the performances of the nearest neighbor classifier. The filtering methods uses score functions that are less expensive to compute than the performance measures and select prototypes based on their individual scores. The wrapper methods is a general, flexible framework for empirically selecting a best model among a finite number of settings, but do not provide an intrinsic framework for learning an optimal model. For the filter methods, strong or complex dependency among prototypes are problematic, as it is difficult to represent such dependency sufficiently by individual scores.

Heuristic combinatorial search such as evolutionary algorithms, a memetic algorithm, an artificial immune model, and tabu search, have been explored to find compact subsets of the training instances [8,16,17,42]. Theoretical approaches for prototype selection shown the upper-bound of generalization error based on VC dimensions, structural risk minimization, and Bayesian analysis [12,15,23]. The looseness of the bounds and the scalability of the algorithms are practical issues in these approaches.

2.2 Approximate Nearest Neighbor Search

Due to the importance of the nearest neighbor classification, other means of pre-processing to reduce the computational requirements for the similarity and distance measures also have been explored in wide areas of study. In recent years, Approximate Nearest Neighbor Search (ANNS) [2] have drawn strong interest and various techniques have been developed for image classification. ANNS uses a family of hash functions to obtain *buckets* to which the prototypes are hashed, and when a test sample is *queried*, its hash value is used to retrieve relevant prototypes from the bucket. The error bounds of the approximation are derived from the probabilities that points closer and further apart than designated distances are hashed to the same and different buckets, respectively.

The ANNS approach yields excellent run-time, typically matching that of the best prototype selection method, while maintaining the full set of prototypes. While the ANNS approach currently receives much more research interest, the prototype selection approach has its merits worth revisiting. As shown in the recent survey [18], the prototype selection methods aimed at reducing the generalization error achieve better accuracy than ANNS. Additionally, the implementation of ANNS requires a class of hash function that can satisfy the conditions for approximation, which is limited to few well-known distance functions, e.g., Hamming, Manhattan, Squared Euclidean, Jaccard, and Arc cos. Finding a suitable LSH family for a domain-specific distance function is generally difficult and is not actively pursued.

This paper focuses on a prototype selection approach which requires only the mutual distance matrix as the input, thus is applicable for problems where ANNS approach and also many prototype selection approaches are not.

2.3 Discriminative Learning and Nearest Neighbor Model

Recently, discriminative learning approaches that have the nearest neighbor classification model embedded in its formulation have been studied in the context of metric learning. Large-margin Nearest Neighbor Classification [38] addresses the problem of distance metric learning as a constrained optimization problem to minimize the violation of the nearest neighbor rule with regards to the affine transformation. The violation is formalized as the margin from the delineating hyperplane and the problem is solved by semi-definite programming. Max-margin Multiclass Nearest Neighbors [25] also learns the metric space based on the minimization of its entropy.

3 Proposed Method

3.1 Preliminaries

Let $\mathcal{X} = \{x_i\}_{i=1}^n$ denote the set of labeled instances and $\mathcal{Y} = \{y_i\}_{i=1}^n$ their class values. Instances are not necessarily represented as vectors, as the nearest neighbor prediction requires only similarities/dissimilarities from the prototypes to a new instance. The nearest neighbor classifier, h, returns the class prediction on a test sample t as follows.

$$h(t) = y_i : \underset{i \in \{1,\ldots,\#(\mathcal{X})\}}{\arg\min} \; d(t, x_i) \tag{1}$$

where $d(\cdot,\cdot)$ denotes the dissimilarity function. A prototype selection algorithm selects a subset $\mathcal{Z} \subseteq \mathcal{X}$ with which the nearest neighbor classifier yields smaller generalization error and run-time.

In order to evaluate a candidate prototype set in the training phase, one must avoid trivial cases where the target instance being classified is also one of the prototypes set, as the prototypes and the training instances originally come from the same set of instances, $\mathcal{X}$. In the following sections, we employ a setup similar to the *leave-one-out* cross validation, such that the target instance is always removed from the prototype set. That is, the prototypes for predicting the class of x_i will be a subset of the remainder of the labeled instances, denote by $\mathcal{Z}_i = \mathcal{X} \setminus \{x_i\}$. For distinction, we refer to an instance as a prototype only when it is used as one of the reference objects for the nearest neighbor rule, in this paper. In turn, we also refer to an instance as a training instance only when it is the subject of prediction.

3.2 Soft-Maximum Function

The *soft maximum* [6] is an approximation of the function $\max(\cdot,\cdot)$ as $\log(\exp(\cdot) + \exp(\cdot))$. The approximation is convex and differentiable, which are desirable for numerical optimization.

A natural extension of the soft maximum for m variables $x_1, \ldots, x_m$ is given by

$$\mathcal{M}(x_1, \ldots, x_m) = \log \sum_{i=1}^{m} \exp(x_i) \tag{2}$$

The base of the exponential, β, is larger than 1 and modified to adjust for the scale of the input values.

3.3 Rank-Adjusted Nearest Neighbor Rule

This section introduces an extension of the nearest neighbor rule in order to parametrize the classification model and the prototype selection problem.

Let x denote a target instance and $\mathcal{Z}$ denote the set of prototypes. We denote by $R(x, x_j; \mathcal{Z})$ the rank of $d(x, x_j)$ in the set $\left\{d(x, x_j)\right\}_{x_j \in \mathcal{Z}}$. The rank takes an integer value from $\{1, \ldots, \#\mathcal{Z}\}$ and a smaller integer is a higher rank. The nearest neighbor rule for a target instance x is as follows.

1. Compute distances $\{d(x, x_j)\}_{x_j \in \mathcal{Z}}$
2. Assign rank $R(x, x_j; \mathcal{Z})$ to all $x_j \in \mathcal{Z}$ based on their distances to x
3. Return the class value of the highest-ranked x_j

We now introduce a positive-value parameter $\alpha(x_j)$ for each prototype to define the adjusted rank R'.

$$R'(x, x_j; \mathcal{Z}) = R(x, x_j; \mathcal{Z}) + \alpha(x_j) \tag{3}$$

The nearest neighbor rule based on the adjusted rank R' generalizes the original rule and produces identical predictions when $\alpha(x_1) = \ldots = \alpha(x_n)$.

Based on the adjusted ranks, the nearest prototype can be passed over for another prototype. We thus refer to $\alpha(x_j)$ as the degradation parameter of the prototype x_j. For brevity, we will use α_j as $\alpha(x_j)$ when it is clear from context. One may correct individual misclassifications by assigning larger degradations to the related prototypes. That is, if x_j and x_k are respectively the highest and the second highest-ranked prototypes for a training case (x, y) where $y_j \neq y$ and $y_k = y$, the misclassification of the basic rule may be corrected by assigning a relatively larger degradation to α_j than α_k.

Our goal in this extension is to learn the degradation parameters that improve the average classification performance over all hold-out cases. Based on their values, we can identify which prototypes are more/less relevant in the prediction of the training set and reduce the prototype set accordingly. The motivation for focusing on ranks rather than distance values is to avoid the issue of scaling. The latter may vary substantially in scale depending on the function, and are more difficult to adjust.

3.4 Approximation of the Nearest Neighbor Rule

Let (x_i, y_i) denote a pair of a hold-out instance and its class label. We denote by $\mathcal{P}$ the subset of prototypes $\mathcal{Z}_i$ with the same label as y_i and $\mathcal{Q}$ that of labels other than y_i. The condition that the nearest neighbor rule correctly predicts y_i is that its nearest neighbor is an element of $\mathcal{P}$. That is,

$$\max\left(\{-d(x_i, p)\}_{p\in\mathcal{P}}\right) > \max\left(\{-d(x_i, q)\}_{q\in\mathcal{Q}}\right) \tag{4}$$

Note that the trivial prediction does not occur because x_i is not an element of $\mathcal{P}$.

In an ideal prototype set, the condition (4) is satisfied over all training cases for all $x_i \in \mathcal{X}$. In general, however, there may not exist such an ideal set. A practical principle for selecting the prototypes is, therefore, to reduce the cases where the conditions are violated as much as possible.

The formulation of such a principal is non-trivial due to the $\max$ function in (4). Alternatively, we rewrite (4) substituting the soft maximum $\mathcal{M}$ and the ranks of $d(x_i, q)$ and $d(x_i, p)$ for $\max$ and the distance values, respectively.

$$\mathcal{M}\left(\{-R(x_i, p; \mathcal{Z}_i)\}_{p\in\mathcal{P}}\right) \geq \mathcal{M}\left(\{-R(x_i, q; \mathcal{Z}_i)\}_{q\in\mathcal{Q}}\right) + 1, \ \ \forall x_i \in \mathcal{X} \tag{5}$$

In essence, (5) describes the same condition as (4) with regards to the ranks of the prototypes of the same and different classes as the hold-out case. The constant 1 is added to the RHS given that R is the rank and takes an integer value.

Substituting (2) to (5), we have

$$\log(\sum_{p\in\mathcal{P}} \exp\left(-R\left(x_i, p; \mathcal{Z}_i\right)\right)) \geq \log(\sum_{q\in\mathcal{Q}} \exp\left(-R\left(x_i, q; \mathcal{Z}_i\right)\right)) + 1$$

Taking the exponential on both sides, we obtain

$$\sum_{p\in\mathcal{P}} \exp\left(-R(x_i, p; \mathcal{Z}_i)\right) - \sum_{q\in\mathcal{Q}} \exp\left(-R\left(x_i, q; \mathcal{Z}_i\right)\right) \geq \rho(x_i) \tag{6}$$

where $\rho(x_i)$ is

$$\rho(x_i) = (\beta - 1) \sum_{q\in\mathcal{Q}} \exp\left(-R\left(x_i, q; \mathcal{Z}_i\right)\right) \tag{7}$$

Since $\mathcal{P} \cup \mathcal{Q} = \mathcal{Z}_i$, the LHS of (6) is rewritten as

$$\sum_{x_j \in \mathcal{Z}_i} \delta(y_i, y_j) \exp\left(-R\left(x_i, x_j; \mathcal{Z}_i\right)\right) \geq \rho(x_i) \tag{8}$$

where

$$\delta(y_i, y_j) = \begin{cases} 1 & \text{if } y_i = y_j \\ -1 & \text{otherwise} \end{cases}$$

Considering (8) as a constraint, we introduce a slack variable ξ_i for each x_i, to represent the violation of the nearest neighbor rule.

$$\sum_{x_j \in \mathcal{Z}_i} \delta(y_i, y_j) \exp\left(-R\left(x_i, x_j; \mathcal{Z}_i\right)\right) \geq \rho(x_i) - \xi_i \tag{9}$$

Each ξ_i takes a non-negative value, and the violation of the nearest neighbor rule over the training set is given as $\sum_{x_i \in \mathcal{X}} \xi_i$.

3.5 Problem Formulation

Combining the rank degradation parameter and the soft-max approximation, we formulate a constrained optimization problem for the parametrized nearest neighbor classification model.

First, we rewrite (9) using the adjusted rank.

$$\sum_{x_j \in \mathcal{Z}_i} \delta(y_i, y_j) \exp\left(-R\left(x_i, x_j; \mathcal{Z}_i\right) - \alpha(x_j)\right) \geq \rho'(x_i) - \xi_i \tag{10}$$

where $\rho'(x_i)$ is the residual term that corresponds to $\rho(x_i)$ in (9), i.e.,

$$\begin{aligned} \rho'(x_i) &= (\beta - 1) \sum_{q \in \mathcal{Q}} \exp\left(-R'\left(x_i, q; \mathcal{Z}_i\right)\right) \\ &= (\beta - 1) \sum_{q \in \mathcal{Q}} \exp\left(-R\left(x_i, q; \mathcal{Z}_i\right) - \alpha(q)\right) \end{aligned} \tag{11}$$

(10) is not a desirable form for a constrained optimization problem because the parameter being trained appears on both sides. That is, $\alpha(x_j)$ on the LHS and $\alpha(q)$ in $\rho'(x_i)$. Alternatively, we replace $\rho'(x_i)$ with $\rho(x_i)$, which is constant for each x_i, and the solution that satisfies (9) will also satisfy (10) since $\rho'(x_i) \leq \rho(x_i)$. For brevity, we denote $\rho(x_i)$ as ρ_i when it is clear in the context.

Secondly, we consider the regularization for the degradation parameters. By definition, the nearest neighbor rule with adjusted ranks behaves the same when all parameters are increased or decreased by the same amount, which is problematic for convergence. The issue can be addressed by the regularization on $\{\alpha(x_j)\}$, to promote less adjustments of ranks given the same amount of violations.

Let λ denote the regularization function on $\{\alpha(x_j)\}$ and C the trade-off coefficient. The constrained optimization problem is formalized as

$$\underset{\alpha_j \geq 0, \xi_i \geq 0}{\arg\min} \ \lambda(\alpha_1, \ldots, \alpha_n) + C \sum_{i=1}^{n} \xi_i \tag{12}$$

subject to

$$\sum_{x_j \in \mathcal{Z}_i} \delta(y_i, y_j) \exp\left(-R\left(x_i, x_j; \mathcal{Z}_i\right) - \alpha_j\right) \geq \rho_i - \xi_i, \ \ \forall x_i \in \mathcal{X}$$

To gain further insight on the above problem, we rewrite the LHS of the constraints as a linear combination of $\mathbf{w} = (w_1, \ldots, w_n)$ and $\mathbf{r}_i = (r_{i1}, \ldots r_{in})$, where $w_j = \exp(\alpha_j)$ and $r_{ij} = \delta(y_i, y_j)\exp(-R(x_i, x_j; \mathcal{Z}_i))$. Using the ℓ-2 norm of $\mathbf{w}$, whose elements monotonically increase with $\{\alpha_j\}$, as the regularizer, we obtain

$$\underset{w_i \geq 0, \xi_i \geq 0}{\arg\min} \ \|\mathbf{w}\|^2 + C \sum_i \xi_i \tag{13}$$

subject to

$$\mathbf{w}^\top \mathbf{r}_i \geq \rho_i - \xi_i, \ \ \forall x_i \in \mathcal{X} \tag{14}$$

In (13), the problem has reduced to a quadratic programming for learning the vector $\mathbf{w}$ of a linear classifier by max-margin principle. In turn, $\mathbf{r}_i$ interprets as a mapping of the hold-out case to a feature vector space. Each feature in $\mathbf{r}_i$ corresponds to a prototype, and characterizes the training instance x_i. Between the exponential definition of $\mathbf{r}_i$ and the regularization on $\mathbf{w}$, the problem induces a sparse representation which highlights the relevant prototypes.

Given the form of (14), the problem is solved by a constrained gradient descent [24]. From the learned degradation values, each prototype is scored by its relevance, i.e., the highest adjusted rank, in the prediction of the hold-out cases,

$$s(x_j) = \alpha_j + \min\left(\{R(x_i, x_j; \mathcal{Z}_i)\}_{x_i \in \mathcal{Z}_j}\right) \tag{15}$$

The prototypes with low scores are removed from the prototype set as they have less effect on the predictions. If the size of the prototype set is fixed, the prototypes are selected by their scores from the highest to the lowest. When the size of the prototypes is not given, one can choose the size of set by cross validation. In the rest of the paper, we refer to the proposed method as Prototype Selection based on Rank-Exponential (REPS) and $\mathbf{r}_i$ as the Exponential Rank representation of x_i.

4 Algorithm

The procedure of prototype selection by REPS is briefly summarized as follows.

- Compute the rank matrix from the input distance matrix
- Compute the parameter values in the constraints
- Solve the Quadratic Programming problem with regards to the degradation parameters
- Eliminate candidates with high degradation parameters

The description of the algorithm with more detail is shown in Algorithm 1.

In order to address large datasets, we used an efficient approximation algorithm for SVM training, shown in Algorithm 2. Algorithm 2 is a modification of the Structural SVM [21] learning algorithm, which obtains the approximate solution to the linear SVM problem using the cutting-plane method. It solves a modified SVM problem where ρ_i is used as the relatively scaled margins of respective prototypes.

With regards to the computational requirements, the space complexity is dominated by that of storing the mutual distance matrix among labeled instances, which is $O(n^2)$. The time complexity is dominated by the computation of ranks, which is $O(n \log n)$ for each instance.

5 Empirical Results

This section presents an empirical study to evaluate REPS using public benchmark datasets.

The empirical results are provided in supplementary materials

Algorithm 1 REPS algorithm

INPUT: Training set $\mathcal{X} = \{x_i\}_{i=1}^n$, regularization parameter C, prototype selection size k
OUTPUT: Prototype set $\mathcal{Z} \subset \mathcal{X}$
METHOD:
Compute the distance matrix $[d(x_i, x_j)]_{n \times n}$ and the ranks $\{R(x_i, x; \mathcal{Z}_i)\}_{x \in \mathcal{Z}_i}$
Compute $r_{ij} = \delta(y_i, y_j) \exp(-R(x_i, x_j; \mathcal{Z}_i))$ and $\rho(x_i)$ (7)
Solve (13) for $\mathbf{w}$
for $j = 1$ to n **do**
  Compute $\alpha_j = \log w_j$ and $s(x_j)$ from (15)
end for
return $\underset{\mathcal{Z} \subset \mathcal{X}: \#\mathcal{Z} = k}{\arg\max} \min(\{s(x_j)\}_{x_j \in \mathcal{Z}})$

Algorithm 2 Cutting-plane Training algorithm

INPUT
OUTPUT
Method
Initialize $\mathcal{W} = \emptyset$
repeat
  $(\mathbf{w}, \xi) \leftarrow \underset{\mathbf{w}, \xi \geq 0}{\arg\min} \frac{1}{2}\mathbf{w}^\top \mathbf{w} + C\xi$ s.t. $\forall c \in \mathcal{W} : \frac{1}{n}\mathbf{w}^\top \sum_{i=1}^n c_i y_i \mathbf{x}_i \geq \frac{1}{n}\sum_{i=1}^n c_i \rho_i - \xi$
  for $i = 1, \ldots, n$ **do**
    $c_i \leftarrow \begin{cases} 1 \text{ if } y_i(\mathbf{w}^\top \mathbf{x}_i < \rho_i) \\ 0 \quad \text{otherwise} \end{cases}$
  end for
  $\mathcal{W} \leftarrow \mathcal{W} \cup \{\mathbf{c}\}$
until $\left\{\frac{1}{n}\sum_{i=1}^n c_i - \frac{1}{n}\sum_{i=1}^n c_i y_i(\mathbf{w}^\top \mathbf{x}_i)\right\} \leq \xi + \epsilon$
return $(\mathbf{w}, \xi)$

5.1 Datasets

For the first part of the experiment, we employed a collection of 34 datasets with vector features from the UCI Machine Learning Repository [4] and the KEEL datasets [1]. The summary of the datasets is shown in Table 4 in Appendix 6. Although some of the benchmarks are not large datasets, a substantial reduction of the prototype set, even when the training set is small, is not irrelevant in practical aspect, as the execution time for the nearest neighbor prediction is a considerable obstacle in its applications.

In the second part of the experiment, we use a collection of 27 datasets from the UCR Time Series Dataset Repository [9]. The input data for time series classification is given as a dissimilarity matrix of Dynamic Time Warping (DTW) [5], which is known to be highly effective in time series classification [13]. We tested window sizes between five and fifty for DTW, and the results were generally consistent among them. The window size is five for all results shown here.

5.2 Baseline Methods

We selected five recent prototype selection methods, which ranked highly among the 42 algorithm reported in a recent survey [18], as baselines for comparative analysis. Class conditional instance selection (CCIS) uses pairwise relations among prototypes to define the selection criteria [27]. Fast condensed nearest neighbor (FCNN) exploits the condensation operation for prototype selection [3]. DROP3 uses a decremental procedure to reduce prototypes and is fast, efficient, and accurate among algorithms of the same approach [39]. Random Mutation Hill Climbing (RMHC) uses stochastic hill climbing search to explore the prototype subset combination space [35]. SSMA conducts a stochastic population search based on the memetic algorithm, and exhibited the best overall

performances among the methods reported in the above survey. All baseline algorithms were executed in the KEEL software [1]. We also report the performance of the nearest neighbor classifier using the full training set as prototypes (NoPS).

The parameter specifications for the baseline methods are summarized in Table 3 in Appendix A. Each method is evaluated using values in the second column and the best result is reported. With regards to the parameters of REPS, the regularization coefficient C was set to 0.001. Its results were robust with regards to the values of C.

5.3 Evaluation

The prototype selection methods are evaluated in two aspects: the accuracy of the nearest neighbor classifier and the compactness of the prototype set, and we use the error rate (ERR) and the selection rate (SLR) for respective measurements. ERR is the ratio of misclassifications with respect to the total number of predictions. SLR is defined as the ratio of the prototype set size with respect to the training set.

Due to the trade-off between the two measures, however, it does not suffice to compare algorithms by single-objectives. Here, we employed three approaches to compare the performance measures with the baselines in multi-objective manners. First, we introduce the concept of the Pareto frontier and non-dominated sorting, commonly used in multi-objective population-based search algorithms [10], to compute the ranks of the algorithms based on two evaluation measures.

Given a problem and a set of algorithms, the Pareto frontier is the group of algorithms to which no other algorithms are better in all evaluation measures. Extending this concept, we can divide the algorithms into separate groups by iteratively finding and removing the Pareto frontier from the working set until no algorithms remain. We then rank each groups of algorithms by how many separate groups of algorithms by which it is dominated. The rank provides a relative measure of performance among the evaluated algorithms, which we will refer to as the Pareto-rank. The exact definition of the Pareto-rank is presented in Appendix C.

Secondly, we evaluated the fixed selection rate (FSR) errors, which is the error rate of the proposed method when its selection rate is forced to the value of each baseline. By forcing the same selection rate, we can utilize the single-measure comparison of the error rates.

Finally, for our sensitivity analysis, we employed the log odds ratio (LOR) of the error and the selection rates. The previous measurements were not suitable as there were no baseline selection rate to which to fix, and the differences between the settings were very small to capture by ranks. Given the error rate of the nearest neighbor algorithm $\mathrm{ERR}_{\mathrm{NoPS}}$, it is defined as

$$\mathrm{LOR} = \log \frac{O(\mathrm{SLR})O(\mathrm{ERR})}{O(\mathrm{ERR} - \mathrm{ERR}_{\mathrm{NoPS}})}$$

where $O(p) = \frac{p}{1-p}$ is the odds given the probability p.

For all measures, a smaller value indicates a better performance. The performances are averaged over 5-fold cross validation for the vector datasets. The default training/test split is used for the time series datasets.

Following the evaluations, we used the Wilcoxon's signed rank test to verify whether the difference between the proposed algorithm and each baseline algorithm is significant. In [11], Wilcoxon's signed ranks test and Friedman-Nemenyi test are recommended for comparison of two classifiers and three or more classifiers over multiple datasets, respectively. Since it is not in our interest to compare and rank among the baseline algorithms, the multiple comparison test result is presented only as a reference. We also note that the Friedman-Nemenyi test is highly conservative, due to the substantial loss of power for handling unreplicated blocked data [28].

5.4 Results

We show large figures and tables in Appendix A and Appendix B, respectively, for better readability of the text.

Table 1 Summary of the Signed Rank Tests

	CCIS	FCNN	SSMA	DROP3	RMHC
PARETO	0.00019	4.6×10^{-6}	0.36	1.1×10^{-6}	0.00043
FSR ERR	7.1×10^{-7}	0.00011	0.39	9.6×10^{-6}	0.11

5.4.1 Vector Data

Fig. 7 in Appendix B illustrates the error and the selection rates of the evaluated algorithms. The x- and y-axes indicate the values of respective measures. Each marker represents the performances of one algorithm on one dataset. Each algorithm is distinguished by a unique color and shape of the marker. The numbers inside the markers identifies the dataset, which are the numbers shown in the first column of Table 7.

Within the same problem (marker number), the markers generally lie from top-left to bottom right, due to the trade-off between the two measures. The shift toward the x-axis indicate the advantage in terms of compactness and the shift towards the y-axis indicate the advantage in accuracy. Note that there is no intrinsic trade-off between the error and selection rates within the same method (marker shape), as the evaluation measures depend on the difficulty of the problems. The balance of the two measures vary depending on the method, e.g., those of REPS and CCIS tend to lean toward higher selection rates.

Table 7 in Appendix A presents more detailed summary of the performances. Each row shows ERR and SLR of each method on one dataset. The selection rate for NoPS, which is always 1, is omitted. The values of algorithms whose Pareto rank is 1 are indicated in bold. It shows that REPS and SSMA are Pareto-rank 1 in the largest numbers of datasets.

Table 2 shows the FSR errors of the proposed algorithm and the error rates of the baselines. The even number columns show the FSR errors of REPS using the same selection rate as the baseline in the next column. The odd numbered columns shows the ERR of the baselines. The selection rates of the baseline algorithms are the same as those shown in through 5^{th}-13^{th} columns of Table 7.

We tested the significance of the differences in the two experimental results, using Wilcoxon's signed rank tests. The error and selection rates from Table 7 are converted into the Pareto ranks of respective algorithms. The ranks of the REPS and each baseline algorithm over all datasets are compared with the alternative hypothesis that the average for REPS is smaller. For comparing the FSR errors, the values from the adjacent, corresponding columns in Table 2 were taken and tested, respectively.

The summary of the tests are shown in Table 1. The first and the second rows show the p-values from the comparisons of the Pareto-ranks and the FSR errors, respectively. In the Pareto-ranks comparison test, the null hypotheses for all baselines but SSMA can be rejected with a high confidence. For the comparison of FSR errors, the null hypotheses for baselines other than SSMA and RMHC can be rejected with a high confidence. SSMA and RMHC were two best algorithms in [18], which indicates that REPS has a significant advantage against many prototype selection algorithms and is competitive with the state-of-the-art.

For the purpose of reference, we include the summary of Friedman rank sum test with post-hoc Nemenyi test on the Pareto rank in the supplementary materials. We note that the Friedman Test with post-hoc Nemenyi test is a very conservative test, and it did not detect significant difference among any pair of methods.

5.4.2 Time Series Data

Table 6 in Appendix A shows the error and the selection rates for the time series datasets. The comparison between REPS and NoPS is shown for this experiment as other baseline methods do not take input of the dissimilarity matrix form. The problems where REPS improves the error rate are indicated by bold. In general, the error rates without prototype selection are generally smaller than those in the previous experiment, which makes the reduction without compromise in the accuracy more difficult.

We show the sensitivity analysis of the logarithm base β of the soft maximum function. We compared the performances of REPS over the same collection of datasets while changing β between 1 and 4.

Fig. 1 summarizes the result of the sensitivity analysis. The y-axis of the box-whisker chart indicates the value of the log odds ratio. It shows that the largest median is achieved with $\beta = 2$, and the performance of REPS is robust with regards to the value of β.

Finally, we present case studies on the examples of the prototypes selected by the proposed method. We focus on the 'ECG200' and 'PLANE' time series. The ECG time series consists of positive (abnormal) and negative (normal) class examples and pose a binary classification problem [29]. The PLANE time series consists of seven classes. For the purpose of the demonstration, classes 1 and 2, whose patterns are mutually the most similar among all pairings, are chosen.

Fig. 3 (1) and (2) show the time series instances of classes 1 and 2 in line plots, respectively. The plots without frames indicate those selected as prototypes and those with dotted frames are the remainders of respective classes in the training set. From these figures, we can see that the overall trends of the two classes are similar, with class 1 having slightly smoother shapes. For both classes, more than half of the training set has been eliminated, and the redundancy of their patterns in relation to the retained prototypes is apparent.

Fig. 2 shows a graph comprised of vertices representing the training examples and edges representing the nearest neighbor relations between the vertices, respectively. The class 2 instances are indicated by the gray background, and the class 1 instances have white backgrounds. The line plots with dotted frames indicate the instances not selected as prototypes. From this figure, we can see that both redundant patterns and potentially erroneous patterns have been thrown out. The class 1 instance on the far right of the graph is potentially erroneous as it neighbors three class 2 instances and has been eliminated.

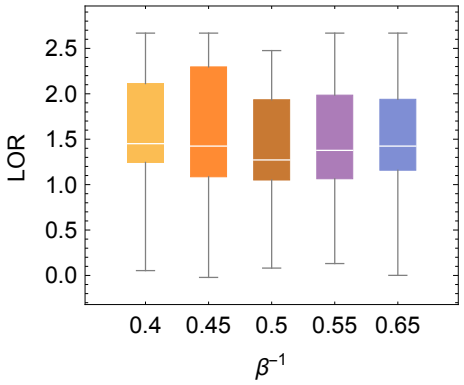

Fig. 1 Sensitivity Analysis for β

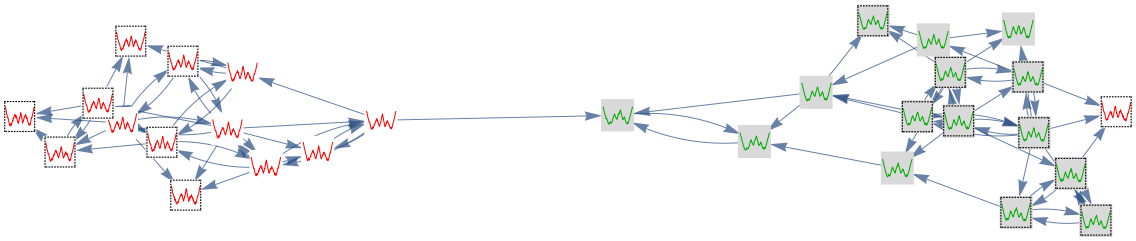

Fig. 2 Prototype Nearest Neighbor Relation Graph (Plane)

In Figs. 4 and 5, the prototypes of positive and negative classes are shown, respectively. The plots without frames indicate those selected as prototypes and those with dotted frames are the remainders of respective classes in the training set. For both classes, a larger number of prototypes are retained compared to the previous examples, as a result of more noise and variations in the patterns.

Fig. 6 shows a graph comprised of vertices that represent the training examples and edges that represent the nearest neighbor relations among the vertices, respectively. In a similar manner to Fig. 2, the gray background indicates positive examples and the white background indicates the negative examples. The vertices with dotted frames are examples not selected as prototypes. As with the previous examples potentially erroneous patterns, neighboring the other class, and few redundant patterns placed among the same classes are eliminated.

These examples demonstrate that the proposed method effectively addresses both the redundant patterns and noisy patterns without exploiting the explicit knowledge of the vector representation space.

6 Conclusion

This paper presented an extension of the nearest neighbor rule based on the adjustment of ranks to parametrize the prototype selection problem and also to approximate the violation of the rule over the training set. As a result, the problem is defined as discriminative learning in a sparse feature space. Our empirical results showed that it is competitive with the state-of-the-art prototype selection algorithm has the advantage over many other existing algorithms in multiple-measure comparisons.

The key intuition for this approach is to avoid intermediate evaluation and to learn the priorities of the prototypes in a direct relation with the performance over the training set. In effect, the proposed approach exploits the sparse relationships among the neighboring instances instead of dense representations such as feature vectors and the distance matrix. Discarding such a large amount of information is justifiable given the successful empirical results of the k-nearest neighbor algorithm in many applications using a small k. When the number of relevant neighbors is limited, a dense representation can be redundant and a sparse representation can induce efficient learning. The empirical results suggests that the sparse representation can maintain the dependence among the prototypes that are stronger and more complex in a reduced set, which is an advantage of the proposed method that and is difficult to achieve using locally-defined criteria and instance-based analyses.

References

1. Alcalà-Fdez, J., Sánchez, L., García, S., del Jesus, M., Ventura, S., Garrell, J., Otero, J., Romero, C., Bacardit, J., Rivas, V., Fernández, J., Herrera, F.: Keel: a software tool to assess evolutionary algorithms for data mining problems. Soft Computing **13**(3), 307–318 (2009).
2. Andoni, A., Indyk, P.: Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. Commun. ACM **51**(1), 117–122 (2008).
3. Angiulli, F.: Fast nearest neighbor condensation for large data sets classification. IEEE Trans. on Knowl. and Data Eng. **19**(11), 1450–1464 (2007).
4. Bache, K., Lichman, M.: UCI Machine Learning Repository. http://archive.ics.uci.edu/ml (2013).
5. Berndt, D.J., Clifford, J.: Using Dynamic Time Warping to Find Patterns in Time Series. In: Proceedings of KDD-94: AAAI Workshop on Knowledge Discovery in Databases, pp. 359–370. Seattle, Washington (1994)
6. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press, New York, NY, USA (2004)
7. Brighton, H., Mellish, C.: Advances in instance selection for instance-based learning algorithms. Data Min. Knowl. Discov. **6**(2), 153–172 (2002).
8. Cano, J.R., Herrera, F., Lozano, M.: Using evolutionary algorithms as instance selection for data reduction in kdd: An experimental study. Trans. Evol. Comp **7**(6), 561–575 (2003).
9. Chen, Y., Keogh, E., Hu, B., Begum, N., Bagnall, A., Mueen, A., Batista, G.: The UCR time series classification archive (2015).
10. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Trans. on Evolutionary Computation **6**, 182–197 (2002)
11. Demšar, J.: Statistical comparisons of classifiers over multiple data sets. J. Mach. Learn. Res. **7**, 1–30 (2006).

12. Devroye, L., Györfi, L., Lugosi, G.: Condensed and edited nearest neighbor rules. In: A Probabilistic Theory of Pattern Recognition, *Stochastic Modelling and Applied Probability*, vol. 31, pp. 303–313. Springer New York (1996).
13. Ding, H., Trajcevski, G., Scheuermann, P., Wang, X., Keogh, E.: Querying and mining of time series data: Experimental comparison of representations and distance measures. Proc. VLDB Endow. **1**(2), 1542–1552 (2008).
14. Fayed, H.A., Atiya, A.F.: A novel template reduction approach for the k-nearest neighbor method. Trans. Neur. Netw. **20**(5), 890–896 (2009).
15. Ferrandiz, S., Boullé, M.: Bayesian instance selection for the nearest neighbor rule. Mach. Learn. **81**(3), 229–256 (2010).
16. Garain, U.: Prototype reduction using an artificial immune model. Pattern Anal. Appl. **11**(3-4), 353–363 (2008).
17. García, S., Cano, J.R., Herrera, F.: A memetic algorithm for evolutionary prototype selection: A scaling up approach. Pattern Recogn. **41**(8), 2693–2709 (2008).
18. Garcia, S., Derrac, J., Cano, J., Herrera, F.: Prototype selection for nearest neighbor classification: Taxonomy and empirical study. IEEE Trans. Pattern Anal. Mach. Intell. **34**(3), 417–435 (2012).
19. García-Pedrajas, N.: Constructing ensembles of classifiers by means of weighted instance selection. Trans. Neur. Netw. **20**(2), 258–277 (2009).
20. García-Pedrajas, N., De Haro-García, A.: Boosting instance selection algorithms. Know.-Based Syst. **67**, 342–360 (2014).
21. Joachims, T., Finley, T., Yu, C.N.J.: Cutting-plane Training of Structural SVMs. Mach. Learn. **77**, 27–59 (2009).
22. Kadous, M.W., Sammut, C.: Classification of Multivariate Time Series and Structured Data Using Constructive Induction. Mach. Learn. **58**(2-3), 179–216 (2005).
23. Karacali, B., Krim, H.: Fast minimization of structural risk by nearest neighbor rule. Trans. Neur. Netw. **14**(1), 127–137 (2003).
24. Kelley, C.: Iterative Methods for Optimization. Frontiers in Applied Mathematics. Society for Industrial and Applied Mathematics (1999).
25. Kontorovich, A., Weiss, R.: Maximum margin multiclass nearest neighbors. In: Proceedings of The 31st International Conference on Machine Learning, vol. 32, pp. 892–900 (2014).
26. Lumini, A., Nanni, L.: A clustering method for automatic biometric template selection. Pattern Recognition **39**(3), 495 – 497 (2006).
27. Marchiori, E.: Class conditional nearest neighbor for large margin instance selection. IEEE Trans. Pattern Anal. Mach. Intell. **32**(2), 364–370 (2010).
28. Nemenyi, P.: Distribution-free multiple comparisons. Ph.D. thesis, Princeton University (1963)
29. Olszewski, R.T.: Generalized feature extraction for structural pattern recognition in time-series data. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, USA (2001). AAI3040489
30. Olvera-López, J.A., Carrasco-Ochoa, J.A., Martínez-Trinidad, J.F.: A new fast prototype selection method based on clustering. Pattern Anal. Appl. **13**(2), 131–141 (2010).
31. Olvera-López, J.A., Carrasco-Ochoa, J.A., Martínez-Trinidad, J.F., Kittler, J.: A review of instance selection methods. Artif. Intell. Rev. **34**(2), 133–143 (2010).
32. Paulevé, L., Jégou, H., Amsaleg, L.: Locality sensitive hashing: A comparison of hash function types and querying mechanisms. Pattern Recogn. Lett. **31**(11), 1348–1358 (2010).
33. Pkalska, E., Duin, R.P.W., Paclík, P.: Prototype selection for dissimilarity-based classifiers. Pattern Recogn. **39**(2), 189–208 (2006).
34. Riquelme, J.C., Aguilar-Ruiz, J.S., Toro, M.: Finding representative patterns with ordered projections. Pattern Recognition **36**(4), 1009 – 1018 (2003).
35. Skalak, D.B.: Prototype and feature selection by sampling and random mutation hill climbing algorithms. In: Proceedings of the Eleventh International Conference on Machine Learning, pp. 293–301. Morgan Kaufmann (1994)
36. Srisawat, A., Phienthrakul, T., Kijsirikul, B.: Sv-knnc: An algorithm for improving the efficiency of k-nearest neighbor. In: Q. Yang, G. Webb (eds.) PRICAI 2006: Trends in Artificial Intelligence, *Lecture Notes in Computer Science*, vol. 4099, pp. 975–979. Springer Berlin Heidelberg (2006).
37. Triguero, I., Derrac, J., Garcia, S., Herrera, F.: A taxonomy and experimental study on prototype generation for nearest neighbor classification. Trans. Sys. Man Cyber Part C **42**(1), 86–100 (2012).
38. Weinberger, K.Q., Saul, L.K.: Distance Metric Learning for Large Margin Nearest Neighbor Classification. J. Mach. Learn. Res. **10**, 207–244 (2009).
39. Wilson, D.R., Martinez, T.R.: Reduction techniques for instance-basedlearning algorithms. Mach. Learn. **38**(3), 257–286 (2000).
40. Wu, X., Kumar, V.: The Top Ten Algorithms in Data Mining, 1st edn. Chapman & Hall/CRC (2009)
41. Xi, X., Keogh, E., Shelton, C., Wei, L., Ratanamahatana, C.A.: Fast Time Series Classification using Numerosity Reduction. In: ICML '06: Proceedings of the 23rd International Conference on Machine Learning, pp. 1033–1040. ACM, New York, NY, USA (2006).
42. Zhang, H., Sun, G.: Optimal reference subset selection for nearest neighbor classification by tabu search. Pattern Recognition **35**(7), 1481 – 1490 (2002).

A Tables

Table 2 Summary of FSR Error Rates

	REPS	CCIS	REPS	FCNN	REPS	SSMA	REPS	DROP3	REPS	RMHC
appendicitis	**0.17**	0.18	**0.17**	0.24	0.22	**0.15**	**0.19**	0.26	0.19	**0.13**
australian	**0.28**	0.44	**0.28**	0.41	**0.34**	0.41	**0.27**	0.39	**0.27**	0.42
balance	**0.12**	0.22	**0.15**	0.30	0.21	**0.12**	**0.14**	0.17	0.15	**0.12**
bands	**0.32**	0.41	**0.32**	0.42	**0.38**	0.42	**0.30**	0.39	**0.36**	0.42
breast	**0.33**	0.38	**0.32**	0.37	0.43	**0.35**	**0.32**	0.33	**0.36**	0.37
bupa	0.39	**0.34**	**0.39**	0.41	0.40	**0.33**	0.41	**0.35**	0.40	**0.33**
cleveland	**0.48**	0.74	**0.47**	0.68	**0.49**	0.55	**0.47**	0.61	**0.46**	0.60
contraceptive	0.54	**0.53**	0.55	**0.53**	0.55	**0.50**	**0.51**	0.52	0.52	**0.50**
crx	**0.23**	0.47	**0.23**	0.44	**0.27**	0.40	**0.26**	0.44	**0.26**	0.42
dermatology	**0.076**	0.61	**0.070**	0.53	**0.076**	0.80	**0.076**	0.80	**0.076**	0.75
ecoli	**0.26**	0.32	0.27	**0.24**	0.27	**0.20**	0.26	**0.26**	0.27	**0.19**
german	**0.32**	0.49	**0.32**	0.42	**0.34**	0.42	**0.31**	0.41	**0.30**	0.39
glass	**0.36**	0.55	0.34	**0.30**	0.48	**0.36**	0.42	**0.40**	0.48	**0.36**
haberman	0.36	**0.36**	0.35	**0.34**	0.34	**0.28**	**0.32**	0.35	0.32	**0.31**
hayes-roth	**0.39**	0.43	0.39	**0.36**	0.44	**0.41**	**0.38**	0.56	**0.45**	0.47
heart	**0.20**	0.40	**0.19**	0.41	**0.21**	0.42	**0.20**	0.44	**0.21**	0.37
housevotes	**0.083**	0.35	**0.086**	0.11	0.11	**0.079**	**0.094**	0.17	0.087	**0.084**
ionosphere	**0.12**	0.36	**0.13**	0.19	0.26	**0.10**	**0.21**	0.25	0.20	**0.11**
iris	**0.073**	0.31	**0.1**	0.41	0.13	**0.040**	0.1	**0.087**	0.087	**0.033**
lymphography	**0.21**	0.28	**0.22**	0.25	0.24	**0.24**	**0.23**	0.34	0.26	**0.24**
mammographic	**0.21**	0.44	**0.21**	0.29	**0.30**	0.34	**0.21**	0.38	**0.23**	0.27
monk-2	0.18	**0.17**	0.16	**0.088**	0.25	**0.037**	0.17	**0.17**	0.19	**0.11**
movement_libras	**0.32**	0.39	0.32	**0.25**	**0.35**	0.36	0.34	**0.29**	**0.39**	0.40
newthyroid	**0.079**	0.33	**0.12**	0.14	**0.11**	0.19	**0.084**	0.15	**0.098**	0.12
pima	**0.29**	0.38	**0.29**	0.35	**0.27**	0.28	**0.26**	0.36	**0.27**	0.29
saheart	**0.37**	0.39	**0.38**	0.41	**0.29**	0.40	**0.31**	0.42	**0.29**	0.36
sonar	**0.20**	0.46	**0.20**	0.35	0.33	**0.32**	**0.21**	0.29	0.33	**0.28**
spectfheart	**0.23**	1.0	**0.23**	0.31	**0.20**	0.21	**0.21**	0.30	**0.20**	0.25
tae	**0.50**	0.59	**0.49**	0.58	**0.51**	0.64	**0.53**	0.54	**0.54**	0.57
vehicle	**0.30**	0.49	**0.34**	0.38	**0.38**	0.45	**0.36**	0.43	**0.38**	0.45
wdbc	**0.044**	0.21	**0.060**	0.077	0.11	**0.10**	**0.072**	0.11	**0.049**	0.077
wine	**0.056**	0.29	**0.051**	0.34	**0.062**	0.29	**0.062**	0.33	**0.062**	0.34
wisconsin	**0.032**	1.0	**0.031**	0.037	0.044	**0.031**	**0.048**	0.062	**0.029**	0.029
yeast	**0.48**	0.55	0.48	**0.48**	0.46	**0.42**	**0.45**	0.48	0.46	**0.42**

Table 3 Parameter Specifications for Baseline Methods

Baseline	Parameters	Distance
FCNN	Neighbor size {1,2,3}	Euclidean
	Pop. size 30	
SSMA	Mut. prob. 0.001	Euclidean
	Cross-over prob. 0.5	
DROP3		Euclidean
RMHC	S/T Ratio	Euclidean

Table 4 Dataset Properties (Vector)

	#Ex.	#Atts.	#Cl.
appendicitis	106	7	2
australian	690	14	2
balance	625	4	3
bands	539	19	2
breast	286	9	2
bupa	345	6	2
cleveland	297	13	5
contraceptive	1473	9	3
crx	690	15	2
dermatology	366	33	6
ecoli	336	7	8
german	1000	20	2
glass	214	9	7
haberman	306	3	2
hayes-roth	160	4	3
heart	270	13	2
housevotes	435	16	2
ionosphere	351	33	2
iris	150	4	3
lymphography	148	18	4
mammographic	961	5	2
monk-2	432	6	2
movement_libras	360	90	15
newthyroid	215	5	3
pima	768	8	2
saheart	462	9	2
sonar	208	60	2
spectfheart	267	44	2
tae	151	5	3
vehicle	846	18	4
wdbc	569	30	2
wine	178	13	3
wisconsin	699	9	2
yeast	1484	8	10

Table 5 Dataset Properties (Time Series)

Data	# Cl.	# Train	# Test	TS Length
50Words	50	450	455	270
Adiac	37	390	391	176
Beef	5	30	30	470
CBF	3	30	900	128
Chlorine	3	467	3840	166
Coffee	2	28	28	286
Diatom	4	16	306	345
ECG200	2	100	100	96
ECG-5Days	2	23	861	136
ECG-CinC	4	40	1380	1639
FaceAll	14	560	1690	131
FaceFour	4	24	88	350
FacesUCR	14	200	2050	131
Fish	7	175	175	463
GunPoint	2	50	150	150
Lightning-2	2	60	61	637
Lightning-7	7	70	73	319
MedicalImages	10	381	760	99
MoteStrain	2	20	1252	84
OliveOil	4	30	30	570
Plane	7	105	105	144
PowerDemand	2	67	1029	24
RobotSurface	2	20	601	70
RobotSurface2	2	27	953	65
Synthetic	6	300	300	60
Trace	4	100	100	275
Two Patterns	4	1000	4000	128

Table 6 Error and Selection Rates (Time Series)

Data	NoPS	ERR	SLR
50words	0.068	0.068	0.96
Adiac	0.42	0.43	0.94
Beef	0.37	0.37	0.97
CBF	0.013	**0.012**	0.97
Chlorine	0.37	0.38	0.99
Coffee	0.036	0.036	0.96
Diatom	0.37	0.38	0.99
ECG200	0.1	0.13	0.88
ECG-CinC	0.37	0.38	0.99
ECG-5Days	0.19	0.30	0.52
FaceAll	0.068	0.068	0.96
FaceFour	0.076	0.076	0.99
FacesUCR	0.027	0.040	0.88
Fish	0.18	**0.17**	0.90
GunPoint	0.060	0.060	0.94
Lighting2	0.23	0.23	0.99
Lighting7	0.24	0.25	0.90
MedicalImages	0.13	0.16	0.90
MoteStrain	0.16	0.17	0.93
OliveOil	0.13	0.13	0.97
Plane	0.0095	**0.019**	0.70
PowerDemand	0.18	0.34	0.72
RobotSurface	0.089	0.098	0.52
RobotSurface2	0.29	**0.28**	0.95
Synthetic	0.080	0.080	0.92
Trace	0.11	0.20	0.52
TwoPatterns	0.32	0.32	0.97

Table 7 Summary of Error and Selection Rates (Vector Data)

ID	Data Name	NoPS	CCIS		FCNN		SSMA		DROP3		RMHC		REPS	
		ERR	ERR	SLR	ERR	SLR	ERR	SLR	ERR	SLR	ERR	SLR	ERR	SLR
1	appendicitis	0.81	**0.18**	**0.031**	0.24	0.32	**0.15**	**0.038**	0.26	0.13	**0.13**	**0.094**	0.20	0.050
2	australian	0.69	0.44	0.030	0.41	0.32	0.41	0.014	0.39	0.11	0.42	0.10	**0.36**	**0.013**
3	balance	0.78	0.22	0.053	0.30	0.33	**0.12**	**0.030**	0.17	0.12	**0.12**	**0.10**	**0.28**	**0.020**
4	bands	0.71	0.41	0.054	0.42	0.48	0.42	0.056	0.39	0.32	0.42	0.099	**0.37**	**0.017**
5	breast	0.63	0.38	0.052	0.37	0.49	0.35	0.024	**0.33**	**0.20**	0.37	0.099	**0.34**	**0.022**
6	bupa	0.61	0.34	0.098	0.41	0.56	**0.33**	**0.057**	0.35	0.30	**0.33**	**0.098**	**0.45**	**0.017**
7	cleveland	0.53	0.74	0.31	0.68	0.61	**0.55**	**0.019**	0.61	0.17	0.60	0.098	**0.48**	**0.045**
8	contraceptive	0.43	0.53	0.23	0.53	0.71	**0.50**	**0.027**	0.52	0.27	0.50	0.099	**0.61**	**0.0064**
9	crx	0.78	0.47	0.031	0.44	0.31	0.40	0.016	0.44	0.12	0.42	0.10	**0.27**	**0.016**
10	dermatology	0.95	**0.61**	**0.047**	0.53	0.12	**0.80**	**0.036**	0.80	0.085	0.75	0.098	**0.073**	**0.048**
11	ecoli	0.79	0.32	0.14	0.24	0.37	**0.20**	**0.059**	0.26	0.16	**0.19**	**0.097**	**0.31**	**0.045**
12	german	0.67	0.49	0.040	0.42	0.50	0.42	0.026	0.41	0.22	0.39	0.10	**0.36**	**0.0092**
13	glass	0.72	0.55	0.14	**0.30**	**0.48**	**0.36**	**0.077**	0.40	0.25	**0.36**	**0.099**	**0.50**	**0.053**
14	haberman	0.65	0.36	0.034	0.34	0.51	**0.28**	**0.020**	0.35	0.20	0.31	0.098	0.33	0.027
15	hayes-roth	0.69	0.43	0.14	**0.36**	**0.49**	**0.41**	**0.066**	0.56	0.24	0.47	0.094	**0.46**	**0.052**
16	heart	0.79	0.40	0.043	0.41	0.37	0.42	0.026	0.44	0.17	0.37	0.097	**0.26**	**0.023**
17	housevotes	0.87	**0.35**	**0.018**	0.11	0.15	**0.079**	**0.026**	0.17	0.061	0.084	0.097	**0.11**	**0.019**
18	ionosphere	0.91	**0.36**	**0.017**	0.19	0.20	**0.10**	**0.033**	0.25	0.088	0.11	0.10	0.24	0.048
19	iris	0.85	**0.31**	**0.032**	0.41	0.13	**0.040**	**0.047**	0.087	0.073	**0.033**	**0.10**	0.053	0.15
20	lymphography	0.64	0.28	0.066	0.25	0.42	**0.24**	**0.057**	0.34	0.18	0.24	0.093	0.25	0.081
21	mammographic	0.80	0.44	0.017	0.29	0.39	**0.34**	**0.011**	0.38	0.14	**0.27**	**0.099**	**0.28**	**0.014**
22	monk-2	0.75	0.17	0.049	0.088	0.077	**0.037**	**0.031**	0.17	0.20	0.11	0.098	0.24	0.19
23	movement_libras	0.76	0.39	0.31	**0.25**	**0.40**	**0.36**	**0.15**	**0.29**	**0.33**	**0.40**	**0.097**	**0.58**	**0.055**
24	newthyroid	0.85	**0.33**	**0.026**	0.14	0.12	**0.19**	**0.030**	0.15	0.12	0.12	0.099	**0.11**	**0.031**
25	pima	0.97	0.38	0.041	0.35	0.46	**0.28**	**0.026**	0.36	0.18	0.29	0.099	**0.32**	**0.010**
26	saheart	0.65	0.39	0.044	0.41	0.51	0.40	0.030	0.42	0.22	0.36	0.099	**0.31**	**0.017**
27	sonar	0.82	**0.46**	**0.046**	0.35	0.30	**0.32**	**0.074**	0.29	0.26	**0.28**	**0.096**	**0.28**	**0.11**
28	spectfheart	0.72	1.0	1.0	0.31	0.43	**0.21**	**0.024**	0.30	0.17	0.25	0.098	**0.20**	**0.024**
29	tae	0.58	0.59	0.16	0.58	0.60	0.64	0.083	**0.54**	**0.20**	0.57	0.099	**0.55**	**0.060**
30	vehicle	0.71	0.49	0.17	**0.38**	**0.49**	**0.45**	**0.065**	**0.43**	**0.24**	0.45	0.099	**0.63**	**0.0092**
31	wdbc	0.94	**0.21**	**0.0075**	0.077	0.12	**0.10**	**0.015**	0.11	0.044	0.077	0.099	**0.049**	**0.082**
32	wine	0.95	**0.29**	**0.025**	0.34	0.13	0.29	0.034	0.33	0.098	0.34	0.098	**0.073**	**0.12**
33	wisconsin	0.96	1.0	1.0	0.037	0.090	**0.031**	**0.0073**	0.062	0.025	**0.029**	**0.099**	0.045	0.012
34	yeast	0.53	0.55	0.22	0.48	0.66	**0.42**	**0.038**	0.48	0.23	0.42	0.099	**0.53**	**0.020**

B Figures

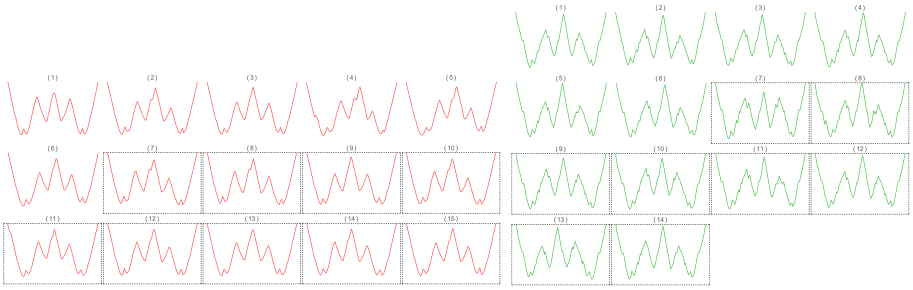

(1) Training set - Class 1 (2) Training set - Class 2

Fig. 3 Prototype Examples (PLANE Time Series)

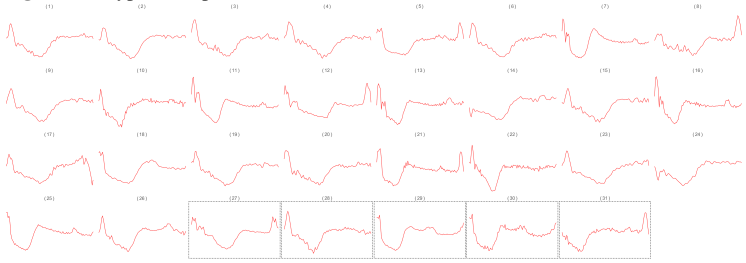

Fig. 4 Prototype Examples (ECG Time Series) - Positive Class

Fig. 5 Prototype Examples (ECG Time Series) - Negative Class

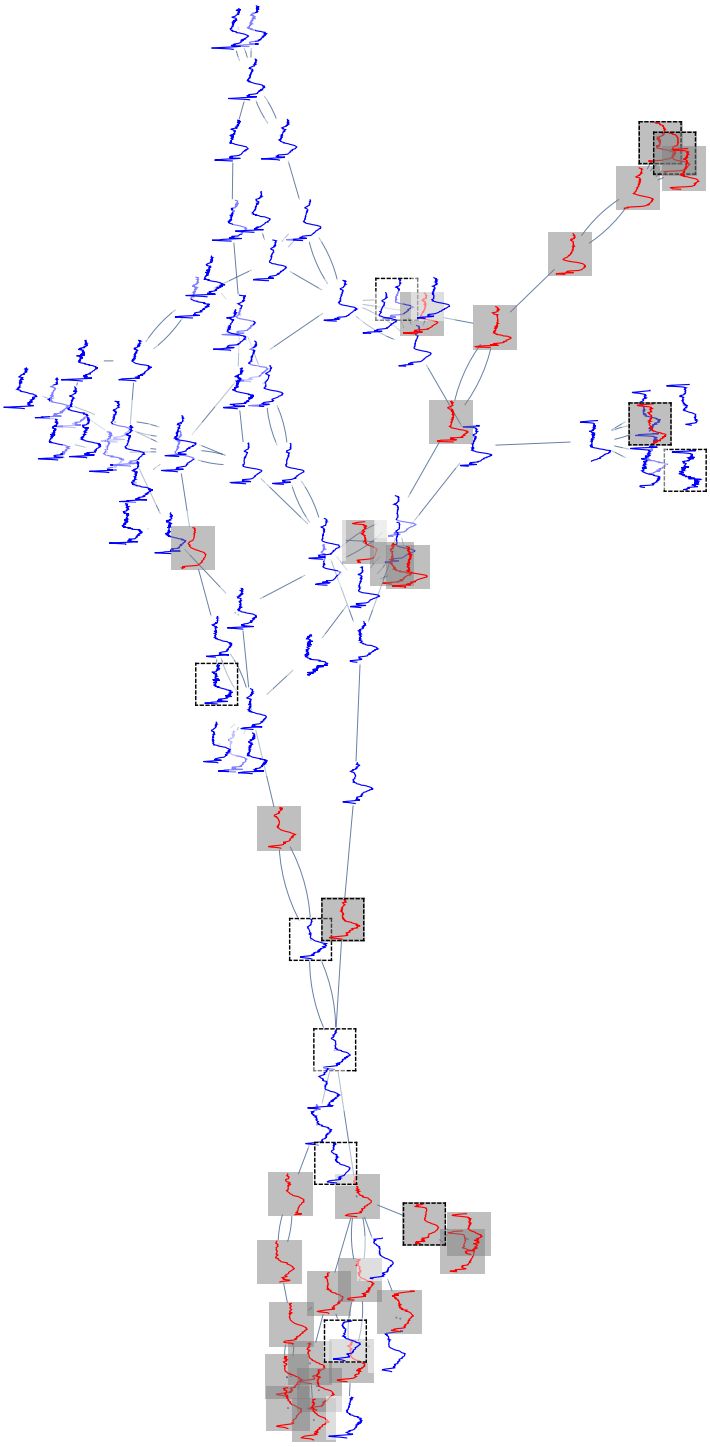

Fig. 6 Prototype Nearest Neighbor Relations Graph (ECG)

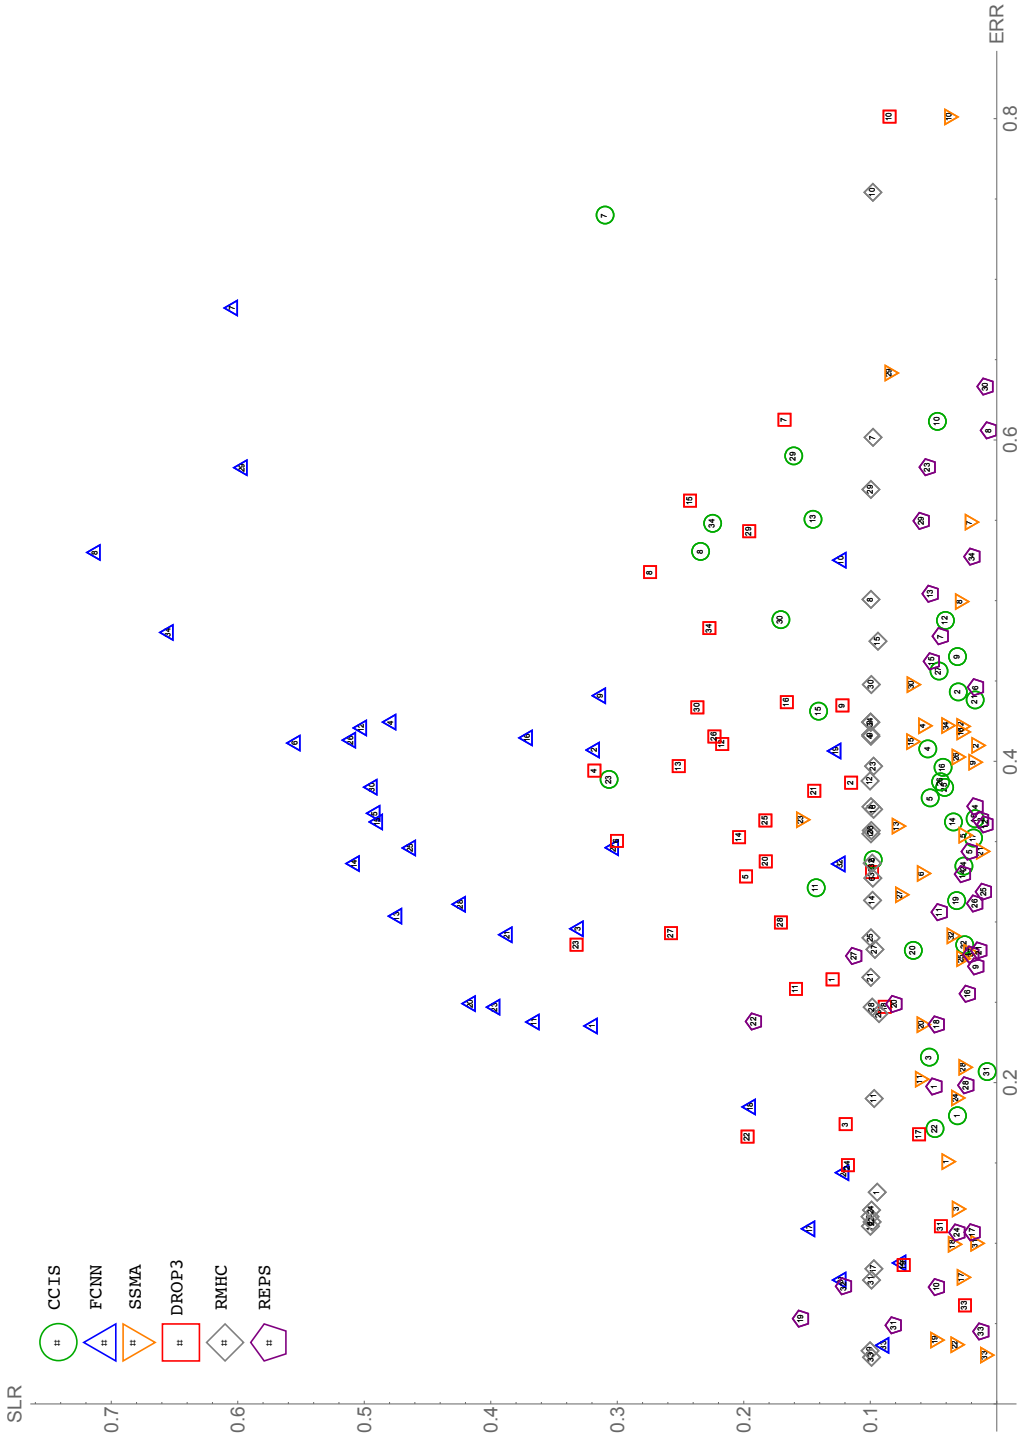

Fig. 7 Error and Selection Rates (Vector Data)

C Definition of Pareto-rank

Let us denote the multi-objective tuple of an algorithm i as $\mathbf{f}(i) = (\text{ERR}_i, \text{SLR}_i)$. The *dominance* between two algorithms i and j is defined as follows: algorithm i is *dominated* by j, if j is better than i in at least one objective and also equal to or better than i in all other objectives.

We denote by $\mathbf{f}_i \prec \mathbf{f}_j$ that $\mathbf{f}_i$ is dominated by $\mathbf{f}_j$, and by $\mathbf{f}_i \not\prec \mathbf{f}_j$ that $\mathbf{f}_i$ is *not dominated* by $\mathbf{f}_j$. The exact definition of dominated and non-dominated relations can be written as follows.

$$\mathbf{f}_i \prec \mathbf{f}_j \text{ if } (\text{ERR}_i \leq \text{ERR}_j \vee \text{SLR}_i \leq \text{SLR}_j) \wedge (\text{ERR}_i < \text{ERR}_j \vee \text{SLR}_i < \text{SLR}_j) \tag{16}$$

$$\mathbf{f}_i \not\prec \mathbf{f}_j \qquad \text{otherwise} \tag{17}$$

Given a set of tuples $\mathcal{P} = \{(\text{ERR}_i, \text{SLR}_i)\}_{i=i}^{p}$, the Pareto frontier is the subset of $\mathcal{P}$ whose members are not dominated by any member of $\mathcal{P}$. Extending the concept of the Pareto frontier, we can divide the algorithms into separate groups by iteratively finding and removing the Pareto frontier from the set until no algorithms remain. We then define the Pareto-rank as the number of groups of algorithms by which it is dominated.

Let $\mathcal{F}_i$ denote the subset of $\mathcal{P}$ with the Pareto-rank i. The exact definition, for $\mathcal{F}_1$ and $\mathcal{F}_i : i > 2$ respectively, is given as follows.

$$\mathcal{F}_1 = \{\mathbf{f} \in \mathcal{P} : \mathbf{f} \not\prec \mathbf{g} \, \forall \, \mathbf{g} \in \mathcal{P}\} \tag{18}$$

$$\mathcal{F}_i = \left\{ \mathbf{f} \in \mathcal{P} : (\mathbf{f} \prec \mathbf{g} \, \forall \, \mathbf{g} \in \mathcal{F}_{i-1}) \wedge \left(\mathbf{f} \not\prec \mathbf{g} \, \forall \, \mathbf{g} \in \mathcal{P} \setminus \bigcup_{h=1}^{i-1} \mathcal{F}_h \right) \right\} \tag{19}$$